arXiv:1409.6155v3 [cs.CV] 5 Oct 2014

1-HKUST: Object Detection in ILSVRC 2014*

Cewu Lu† Hao Chen*‡ Qifeng Chen*§ Hei Law*† Yao Xiao*† Chi-Keung Tang†
†Hong Kong University of Science and Technology
‡The Chinese University of Hong Kong
§Stanford University

Abstract

The Imagenet Large Scale Visual Recognition Challenge (ILSVRC) is the one of the most important big data challenges to date. We participated in the object detection track of ILSVRC 2014 and received the fourth place among the 38 teams. We introduce in our object detection system a number of novel techniques in localization and recognition. For localization, initial candidate proposals are generated using selective search, and a novel bounding boxes regression method is used for better object localization. For recognition, to represent a candidate proposal, we adopt three features, namely, RCNN feature, IFV feature, and DPM feature. Given these features, category-specific combination functions are learned to improve the object recognition accuracy. In addition, object context in the form of background priors and object interaction priors are learned and applied in our system. Our ILSVRC 2014 results are reported alongside with the results of other participating teams.

Keywords: Object Detection, Deep learning and ILSVRC.

1. Introduction

We present our system for the ILSVRC 2014 competition. The big data challenge has evolved over the past few years as one for the most important forums for researchers to exchange their ideas, benchmark their systems, and push breakthrough in categorical object recognition in large-scale image classification and object detection.

The 1-HKUST team made its debut in this year's competition and we focused on the object detection track. The object detection problem can be divided into two subproblems, namely, localization and recognition. Localization solves the "where" problem, while recognition solves the "what" problem. That is, we locate where the objects are, and then recognize which object categories the detected objects should belong to.

We made technical contributions on both localization and recognition. For localization, we exploit regression on bounding box using deep learning based on selective search outputs. For recognition, we focus on integrating the state-of-the-art computer vision techniques to build a more powerful category-specific category predictor. In addition, object background priors are also considered.

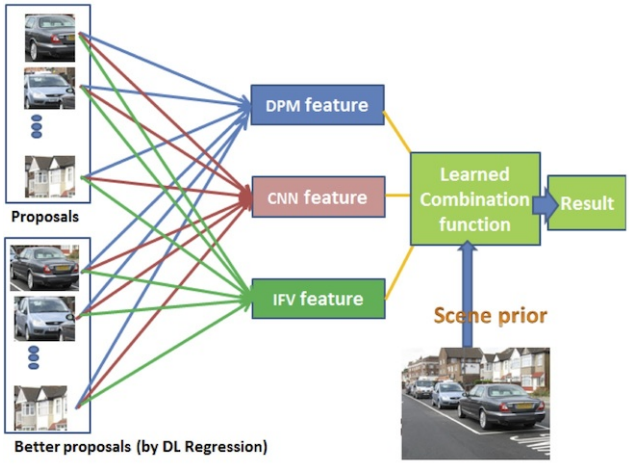

Figure 1. Our framework.

2. Framework

Figure 1 gives the overview of our system which summarizes our contributions in object localization and recognition.

2.1. Localization

We first extract candidate objectness proposals using selective search. As widely known, the output bounding boxes are almost never perfect and fail to coincide the ground-truth object boxes with a high overlap rate (e.g. 80%). To cope with this problem, we learn a regressor using deep learning.

*indicates equal contribution, all the authors are united under **1-HKUST** although some of their current affiliations are elsewhere.

2.2. Recognition

Given a set of candidate proposals in hand, we extract different types of feature representation for recognition. We adopt three types of feature, namely, CNN feature, DPM feather and IFV feature, to measure the given candidate proposals.

For CNN feature, we first train the CNN model similar with CaffeNet (refer to [3] for architecture details), and the outputs of the Fc6 layer are extracted as the CNN features. We apply the SVM training to obtain 200 object category classifiers, as similarly done in RCNN [2]. For DPM feature [1] we also train 200 DPM models. For IFV feature [5], we make use of the fast IFV feature extraction solution [4] to compute at a rate of 20 seconds per image. We also train 200 SVM category models as similarly done for the above two features. After obtaining 200 CNN scores, 200 DPM scores, and 200 IFV scores, these scores are concatenated into a 600-dimensional feature vector. Finally, we train a 200-class SVM model on these features.

2.3. Background Prior

Objects occur in context and are part of the scene. Background scene understanding can definitely benefit object detection. The background can reject (or re-score) unreasonable objects. For example, a yacht does not appear in an indoor environment with high probability. In our implementation, we train a presence prior model (PPM) under the CNN framework on the object detection data of ILSVRC 2014. Rather than producing a single label per image, this method outputs multiple labels for an image. Thus, false predictions can be removed if the prediction score based on our trained presence prior falls below a confidence threshold. Our experimental results demonstrate that the presence prior could help to filter false predictions with more context information being considered.

3. Results

We discuss the performance of our entries in ILSVRC 2014. In the object detection track, there are two subtracks: with and without extra training data. Our results were achieved *without* extra training data. We were ranked fourth in terms of number of winning categories. Table 1 tabulates the top winners and we refer readers to [6] or the official website of ILSVRC 2014 for complete standings. Our mAP is 0.289. By analyzing the per-class results[1], we found that 1-HKUST is still ranked fourth among all the teams using and without using extra training data. Table 2 shows that using extra training data gives a clear advantage. Sample visual results are demonstrated in Figure 2. Surprisingly, a number of difficult cases for human detection such as the lizard in Figure 3 can be reliably detected by our system.

Team name	Number of object categories won	mAP
NUS	106	0.372
MSRA	45	0.351
UvA-Euvision	21	0.320
1-HKUST	**18**	**0.289**
Southeast-CASIA	4	0.304
CASIA-CRIPAC-2	0	0.286

Table 1. Number of object categories won without extra training data.

Team name	Number of object categories won
GoogLeNet	138
CUHK-DeepID-Net	28
Deep-Insight	27
1-HKUST (run 2)	**3**
Berkeley-Vision	1
NUS	1
UvA-Euvision	1
MSRA-Visual-Computing	0
MPG-UT	0
ORANGE-BUPT	0
Trimps-Soushen	0
MIL	0
Southeast-CASIA	0
CASIA-CRIPAC-2	0

Table 2. Number of object categories won with and without extra training data. 1-HKUST did not use extra training data.

Unlike other participating teams, 1-HKUST had very limited computing budget and resources in our training and experiments: one 24-core server PC (Dell PowerEdge R720 2 x 12C CPU, 128GB RDIMM memory and one NVIDIA GRID K1 GPU), and one 6-core PC (Dell Alienware Aurora 4.1Ghz, 6C CPU, 32GB DDR2 memory and one NVIDIA GeForce GTX 690 GPU). Due to limited computing resources, the parameter tuning might not have been optimized, and we strongly believe that our framework could achieve a better mAP rating if more computing resources available and careful optimization tuning.

[1] http://image-net.org/challenges/LSVRC/2014/results/ilsvrc2014_perclass_results.zip

References

[1] R. D. Felzenszwalb P, McAllester D. A discriminatively trained, multiscale, deformable part model. *PAMI*, 2010. 2

[2] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and se-

Figure 2. Sample object detection results.

mantic segmentation. *CVPR*, 2014. 2

[3] Y. Jia. Caffe: An open source convolutional architecture for fast feature embedding. http://caffe.berkeleyvision.org/, 2013. 2

[4] C. G. M. S. Koen E. A. van de Sande and A. W. M. Smeulders. Fisher and vlad with flair. In *CVPR*, 2014. 2

[5] M. T. Perronnin F, Snchez J. Improving the fisher kernel for large-scale image classification. In *CVPR*, 2010. 2

[6] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge, 2014. 2

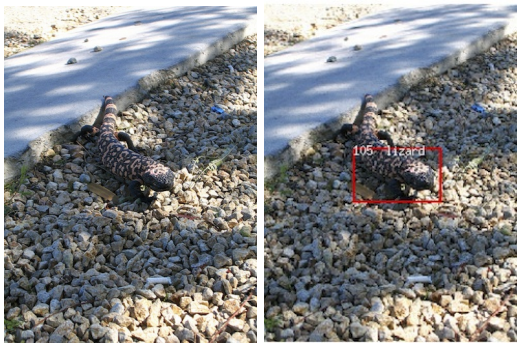

(a) (b)

Figure 3. (a) is a input image, (b) is our detection result. Some people found it difficult to recognize a lizard on pebbles.